

REMARKS

I. Introduction

In response to the Office Action dated February 1, 2006, claims 2-4, 19 and 28 have been amended. Claims 1-40 remain in the application. Re-examination and re-consideration of the application, as amended, is requested.

II. Claim Amendments

Applicants' attorney has made amendments to claims 2-4, 19 and 28 as indicated above. These amendments were made solely for the purpose of clarifying the language of the claims, and were not required for purposes of patentability.

III. Cited References and the Present Invention

A. The Gibson Reference

Gibson et al., "File Server Scaling with Network-Attached Secure Disks", ACM June 1997, pages 272-284, discloses by providing direct data transfer between storage and client, network-attached storage devices have the potential to improve scalability for existing distributed file systems (by removing the server as a bottleneck) and bandwidth for new parallel and distributed file systems (through network striping and more efficient data paths). Together, these advantages influence a large enough fraction of the storage market to make commodity network-attached storage feasible. Realizing the technology's full potential requires careful consideration across a wide range of file system, networking and security issues. The paper contrasts two network-attached storage architectures-(1) Networked SCSI disks (NetSCSI) are network-attached storage devices with minimal changes from the familiar SCSI interface, while (2) Network-Attached Secure Disks (NASD) are drives that support independent client access to drive object services. To estimate the potential performance benefits of these architectures, an analytic model is developed and trace-driven replay experiments are performed based on AFS and NFS traces. The results suggest that NetSCSI can reduce file server load during a burst of NFS or AFS activity by about 30% and with the NASD architecture, server load (during burst activity) can be reduced by a factor of up to five for AFS and up to ten for NFS.

B. The Present Invention

The present invention provides a digital data processing system with improved access to information stored on a peripheral device. The system has a plurality of nodes, a peripheral device,

a file system and a bypass mechanism. A first node (e.g., a client node) is connected to a second node (e.g., a server node) over a first communications pathway (e.g., a network). The second node is itself connected to a peripheral device (e.g., a disk drive) over a second communications pathway. The first node, too, is connected to the peripheral device over a third communications pathway. The file system, executing on the first and second nodes, is capable of responding to access requests generated by the first node for transferring data between that node and the peripheral device, via the second node and via the first and second communications pathways. The file system also maintains administrative information pertaining to storage on the peripheral device of data designated by such requests. That information includes, for example, physical storage location mappings for files and other data stored on the peripheral device. The bypass mechanism, which executes on at least the first node, intercedes in the response to at least selected input/output, or access, requests generated by that node. The bypass transfers data designated by such requests between the first node and the peripheral device over the third communications pathway, in lieu of transferring that data via the second node and the first and second communications pathways. Such transfers by the bypass, however, are made using the administrative information maintained by the file system relating to storage of such data on the peripheral device.

IV. Office Action Prior Art Rejections

In paragraphs (6)-(7), the Office Action rejected claims 1-40 under 35 U.S.C. 102(a) as being anticipated by Gibson.

Respecting claims 1, 2, 3, 4, 16, 19 and 20, the Office Action asserts that Gibson teaches a digital data processing system with improved access to information stored on a peripheral device. The Office Action recites the following passages from Gibson to support the assertion.

“By providing direct data transfer between storage and client, network-attached storage devices have the potential to improve scalability for existing distributed file systems (by removing the server as a bottleneck)...The disk drive industry anticipated the marginal cost for on-disk Fibre Channel interfaces, relative to the common single-ended SCSI interface...**Seagate’s Barracuda FC** is already providing packetized SCSI through **Fibre Channel network ports to directly attached hosts**...Instead we focus on selecting **a command interface that reduces the number of client-storage interactions that must be relayed through the file manager**...Common, data intensive operations, such as reads and writes, offloading more of the file manager’s work without integrating file system policy into the disk...while policy decisions are made in the file manager...Authorization, in the form of a time-limited capability applicable to the **file’s map and contents, should be provided by the file manager** to protect higher-level file systems’ control over storage access policy. While a single drive

object will suffice to represent a simple client file, multiple objects **may be logically linked by the file system** into one client file....As an example of a possible NASD access sequence, consider a file read operation depicted in Figure 3. Before issuing its first read of a file, the client authenticates itself with the file manager and requests access to the file. If access is granted, the client receives the network location of the NASD drive containing the object and a time-limited capability to access the object and for establishing a secure communication channel with the drive. After this point, the client may directly request access to data on NASD drives, using the appropriate capability.” [emphasis added by Office Action]

The Office Action asserts that the preceding text excerpts indicate, inter alia, that “the client only needs initial file mapping authorization, administrative information, i.e., file meta data from the server/file manager” and “after that all read, write or other file system operations go directly to the storage device”. The Office Action further asserts that “a command interface (i.e., filter driver) in the client/first node reduces/filters file system requests that need to be relayed through the server/file manager and hence bypassing the server” and “the administrative information at the server/file manager includes file mappings, i.e. data containing actual physical locations in the storage disk”.

Respecting claims 5-15, 17-18, 20-27 and 29-40, the Office Action merely asserts that the limitations of these claims are addressed in the rejected claim above and refers to those explanation and pages 272-277 of Gibson. Applicants note that the Office Action does not designate the particular part of Gibson relied upon as required under 37 CFR §1.104(c)(2) and Applicants cannot determine where each of the elements of claims 5-15, 17-18, 20-27 and 29-40 is taught by Gibson as required to support a proper §102 rejection.

Applicants respectfully traverse these rejections, because nowhere does Gibson teach or suggest a bypass, executing on at least the first node, for interceding in response to an access request applied thereby to the file system, by transferring data designated by that request ... in accord with administrative information maintained by the file system pertaining to storage of that data as presently claimed, e.g. in claim 1. The bypassing (or intercepting, e.g. as used in claim 19) of an access request avoids directing the access request to the file manager in the usual manner while still using the administrative information (e.g., file maps) from the file manager as presently claimed in each of the independent claims 1-4, 16, 19 and 28.

In general, Gibson outlines a taxonomy of Networked-Attached Storage, wherein:

This paper contrasts two network-attached storage architectures – (1) Networked SCSI disks (NetSCSI) are networked-attached storage devices with minimal changes from the familiar SCSI interface, while (2) Networked-Attached Secure Disks

(NASD) are drives that support independent client access to drive object services.
See Abstract

The Office Action appears to primarily rely on the teachings of Gibson that are directed to the NASD architecture in making the rejections. Almost all the substantive teaching to support the rejection is taken from section 3.4 of Gibson directed to “Case 3: Network-attached Secure Disks (NASD)”. However, the Office Action fails to recognize that operation of the NASD architecture as taught by Gibson requires that “because clients directly request access to data in their files, a NASD drive must have sufficient metadata to map and authorize the request to disk sectors”. See page 275, section 3.4, second paragraph, lines 9-11. Thus, administrative information (e.g. file maps) as taught by Gibson resides on the NASD drive and not with the file manager as presently claimed. Moreover, Applicants submit that Gibson teaches away from using administration information from the file manager as presently claimed by stating that a NASD drive “must have” the metadata to map and authorize the request to disk to support clients directly requesting access to the data.

Applicants further note that the contrasted architecture of NetSCSI described by Gibson also does not teach the present invention as claimed. In this case, Gibson teaches that the file manager “processes the request from a client in the usual way.” See page 275, Figure 2 inset. Thus, although the file manager is eliminated from the data path, the file manager still receives and process the access requests from the clients to the peripheral device. Accordingly, here Gibson does not teach a bypass interceding in response to an access request applied thereby to the file system as presently claimed. Moreover, Applicants submit that Gibson teaches away from such a bypass by teaching the desire (with respect to NetSCSI) to “retain as much as possible of SCSI, the current dominant mid- and high-level storage device protocol” and defines “NetSCSI is a network-attached storage architecture that makes minimal changes to the hardware and software of SCSI disks. See page 274, Section 3.2, first paragraph.

A proper §102 rejection requires that each and every element of the claimed invention is taught by the cited reference. However, Gibson does not teach or suggest the bypassing (or intercepting) of an access request to a file manager while still using the administrative information (e.g., file maps) from the file manager as presently claimed in each of the independent claims 1-4, 16, 19 and 28. Accordingly, Applicants respectfully submit the present rejection of the independent claims under 35 U.S.C. § 102(a) is improper. Withdrawal of the rejection is respectfully solicited.

Further, Applicants submit that dependent claims 5-15, 17, 18, 20-27 and 29-40 are allowable over Gibson in the same manner because they recite all the limitations of their respective independent claims. In addition, dependent claims 5-15, 17, 18, 20-27 and 29-40 recite further novel

features not taught or suggested by Gibson. Accordingly, withdrawal of the rejections and allowance of dependent claims 5-15, 17, 18, 20-27 and 29-40 is respectfully solicited.

V. Conclusion

In view of the foregoing, it is submitted that this application is now in good order for allowance and such allowance is respectfully solicited. Should the Examiner believe minor matters still remain that can be resolved in a telephone interview, the Examiner is urged to call Applicants' undersigned attorney.

Respectfully submitted,

Canady & Lortz LLP
2540 Huntington Drive, Suite 205
San Marino, California 91108
(626) 292-7743 x101
Fax: (909) 494-4441

By:  _____

Name: Bradley K. Lortz
Reg. No.: 45,472

Date: 6/1/06 _____